

Reliable Extraction of the Camera Motion using Constraints on the Epipole

Jonathan Lawn and Roberto Cipolla

Department of Engineering, University of Cambridge

Abstract. An accurate estimate of the epipole (direction of camera translation) is necessary if image motion is to be decomposed into rotational and translational components, which give the camera rotation and feature depths respectively. In this paper we introduce the Linearised Subspace Method to find direct constraints on the epipole which are independent of camera rotation and scene structure. We present methods to compute reliable constraints and their uncertainties from image motion. We show how erroneous constraints due to errors in tracking can be rejected and how the valid constraints should be combined to form accurate estimates of the direction of translation. Experimental results show these methods lead to improvements in the recovery of camera motion and that the uncertainty estimates are accurate and useful in detecting degenerate scene structure or camera motions.

1 Introduction

In the structure from motion problem image motion is used to recover unknown camera motion and scene structure. The Fundamental Matrix [2] is commonly used to encode the camera motion (up to a projective transformation) and the rigidity or epipolar constraint. It is independent of scene structure; can be estimated from point correspondences and has successfully been exploited in segmentation and outlier rejection algorithms [14]. In principle the Fundamental matrix can be decomposed to find the epipole (direction of translation) and hence the camera motion and scene structure up to a projective transformation. In practice this is very ill-conditioned and leads to inaccurate estimates of the epipole.

An accurate estimate of the epipole is, however, an essential component of the structure-from-motion problem if motion and scene structure is to be recovered. Finding the epipole accurately removes much of the uncertainty from the calculation of the rotational component of visual motion, and therefore also improves the calculation of the feature depths [6, 13], and the rejection of outliers caused by spurious or mis-tracked features.

This paper proposes a novel method – the Linearised Subspace Method– of finding linear constraints on the epipole from small regions of a full perspective image. Perturbation analysis is used to produce estimates of the uncertainty in these constraints and reliable constraints are combined to produce an accurate estimate of the epipole even in the presence of outliers and degenerate scene structure and motion. Experimental results are presented.

2 Constraints on the Epipole

If the parameters of the motion of the camera between frames $k - 1$ and k (in the co-ordinate frame of k) are a rotation, \mathbf{R} , and a translation, \mathbf{T} , then the motion of each point is given by the rigid body transformation:

$$\mathbf{X}'_i = \mathbf{R}(\mathbf{X}_i + \mathbf{T}) \quad (1)$$

where \mathbf{X}'_i and \mathbf{X}_i are the vectors to the stationary feature point, i , in the coordinate frames of the frames $k - 1$ and k respectively. Points in the images are represented by the 3-dimensional vectors \mathbf{P}'_i and \mathbf{P}_i , where $\mathbf{X}_i = r_i \mathbf{P}_i$ and r_i is the unknown scalar depth.

The Subspace method [7] for obtaining a constraint on the position of the epipole is based on cancelling the rotational component of motion in a similar way to motion parallax [12]. A discrete time/view version of the Subspace method can be obtained by rearranging the rigid motion equation (1) into the product relationship below.

$$\mathbf{T} \cdot (\mathbf{P}'_i \times \mathbf{P}_i) = \rho_i \mathbf{T} \cdot (\mathbf{P}'_i \times (R^{-1} \mathbf{P}'_i)) \quad (2)$$

Approximating $\rho_i = \frac{r'_i}{r_i}$ as a constant, $\bar{\rho}$, which is intrinsic to the continuous time formulation used originally [7], and taking a weighted sum across a set of points:

$$\mathbf{T} \cdot \sum_i c_i (\mathbf{P}'_i \times \mathbf{P}_i) \approx \bar{\rho} \mathbf{T} \cdot \sum_i c_i (\mathbf{P}'_i \times (R^{-1} \mathbf{P}'_i)) \quad (3)$$

The weights, c_i , can be chosen to cancel the second order terms of \mathbf{P}'_i by making

$$\sum_i c_i \mathbf{P}'_i \mathbf{P}'_i^T = 0 \quad (4)$$

and then the right hand side of Equation 3 will equal zero, and therefore

$$\mathbf{T} \cdot \sum_i c_i (\mathbf{P}'_i \times \mathbf{P}_i) \approx 0 \quad (5)$$

for any camera rotation, R . The linear sum cancels the rotational component of the visual motion, just as in Affine Motion Parallax [1, 10], giving a constraint on the direction of translation. The constraint on the weights, c_i , in Equation 4 can be rewritten as

$$\begin{bmatrix} p'_{xi} p'_{xi} & & \\ \cdots p'_{yi} p'_{xi} & \cdots & \\ p'_{zi} p'_{xi} & & \\ p'_{yi} p'_{yi} & & \\ \cdots p'_{zi} p'_{yi} & \cdots & \\ p'_{zi} p'_{zi} & & \end{bmatrix} \begin{pmatrix} c_1 \\ c_2 \\ \vdots \end{pmatrix} = [\phi_1 \phi_2 \cdots] \mathbf{c} = \Phi \mathbf{c} = \mathbf{0} \quad (6)$$

where $\mathbf{P}'_i = [p'_{xi} \ p'_{yi} \ p'_{zi}]^T$, and Φ is a $6 \times n$ matrix with columns, ϕ_i made up of the 6 independent elements of the matrix $\mathbf{P}'_i \mathbf{P}'_i^T$. Given enough points ($n > 6$), a subspace of possible weights, \mathbf{C} , can be found.

$$\Phi [c_1 \ c_2 \ \dots] = \Phi \mathbf{C} = 0 \quad (7)$$

This space is invariant to affine deformations of the image [7]. It is therefore possible to write (from Equation 5)

$$\mathbf{T}^T \mathbf{M} \mathbf{C} = 0^T \quad (8)$$

where the $3 \times n$ matrix \mathbf{M} is defined as

$$\mathbf{M} = [(\mathbf{P}'_1 \times \mathbf{P}_1) (\mathbf{P}'_2 \times \mathbf{P}_2) \dots] \quad (9)$$

Since \mathbf{M} and \mathbf{C} can be computed from measurements in the image, each column of the matrix $\mathbf{M} \mathbf{C}$ provides a constraint on the direction of translation and the epipole. These constraints are in fact equivalent to epipolar lines.

2.1 The Linearised Subspace Method

The accuracy and sensitivity to noise of the Subspace algorithm can be improved by looking at constraints from *small* regions of the image. If small image regions are used, then the equation for each weight vector, c_a (Equation 4) is approximately linear [11].

$$\sum_i \hat{c}_i \mathbf{P}'_i = \mathbf{0} \quad (10)$$

(Small here means in a region of radius less than 0.1 radians approximately, so that the first order variations from the mean image position are significantly greater than the second order terms.) Only four points are now needed to extract a constraint instead of seven using the Heeger and Jepson Subspace method. This has advantages for outlier rejection (see Section 5). The linearisation may also improve the stability of the solution to noise on the image motion measurements.

Each region should only produce one constraint because under the weak perspective assumption for the small field of view in each neighbourhood the constraints will have the same direction. In the original algorithm [7], an arbitrary weight vector was chosen. Here we choose the constraint with the least uncertainty (as computed below). This takes into account image localisation errors and degeneracy in the structure being viewed.

Though each constraint, \hat{c}_a , (generated from four or more points) must come from a small region, different constraints can be generated from different small image regions in a full perspective image and combined to estimate the epipole. *This method is very similar to the Affine Motion Parallax algorithm [10], generalised for any number of feature points (≥ 4)*. It has the same geometric intuitiveness: combine the visual motions of a number of features from a small

image region in a linear sum, dependent on their positions only, to cancel the rotational component of the motion.

Figure 1(a) shows the image motion from a sequence of views of Trinity College. The image displacement vectors were automatically obtained by tracking “corner” features [5]. The Linearised Subspace method was used to compute constraints on the epipole by looking at 14 image regions. These are displayed in 1(b).

3 Uncertainty Analysis

For real, noisy data the epipole constraints will not be exact, and therefore constraints from more than two image regions will not intersect. A method is needed that will provide an estimate of the accuracy of each constraint. This will allow *excessively uncertain* constraints to be rejected; *optimal combination* of the good constraints to provide the least uncertain result possible; *degeneracy* of the structure or motion to be detected by measuring the significance of components quantitatively, and *outliers* (due to independent motion and incorrect visual motion measurements) to be detected.

The feature measurement errors on the image plane (which may be an output of the feature detector, or be considered constant and determined in advance) will determine the uncertainty in the constraint. The propagation of the image measurement errors to compute the uncertainty in the constraint is in fact directly linked to the camera motion and 3D structure of the points being viewed.

The approach to propagating the errors through the calculations is to assume that each vector or scalar (eg. \mathbf{b}) has a small additive gaussian uncertainty which can be represented by a covariance matrix ($\mathcal{V}[\mathbf{b}]$) [8]. These error predictions are then propagated through the calculations by linearising around each result algebraically. Differentiating about each input variable (\mathbf{a}) gives the Jacobian ($J_{\mathbf{a}}(\mathbf{b})$) and this shows how the input covariances contribute to the output covariance.

The constraint vector, \mathbf{Mc} , will have two elements to its uncertainty: $\delta\mathbf{M}\mathbf{c}$ and $\mathbf{M}\delta\mathbf{c}$. The contribution of both can be calculated. However, we would expect the “velocity component” ($\delta\mathbf{M}$) to be more uncertain than the component dependent on image position. We therefore expect the first term, $\delta\mathbf{M}\mathbf{c}$, to be the more significant. Experiments validate this [11]. We therefore calculate the variance on the constraint to be:

$$\mathcal{V}[\mathbf{Mc}] \approx \sum_i c_i^2 \mathcal{V}[\mathbf{m}_i] = \sum_i c_i^2 \mathcal{V}[\mathbf{P}_i \times \mathbf{P}'_i] \quad (11)$$

$$= \sum_i c_i^2 \left(\mathbf{P}'_{i\times} \mathcal{V}[\mathbf{P}_i] \mathbf{P}'_{i\times}^T + \mathbf{P}_{i\times} \mathcal{V}[\mathbf{P}'_i] \mathbf{P}_{i\times}^T \right) \quad (12)$$

where \mathbf{a}_{\times} is the skew-symmetric matrix that produces the vector product with \mathbf{a} . Figures 1(c) and (d) show two constraints and their respective uncertainties. The uncertainty on the constraints depends on 3D structure of the points as well as the image localisation errors. Nearly planar scene structure results in a very uncertain constraint (Figure 1(d)).

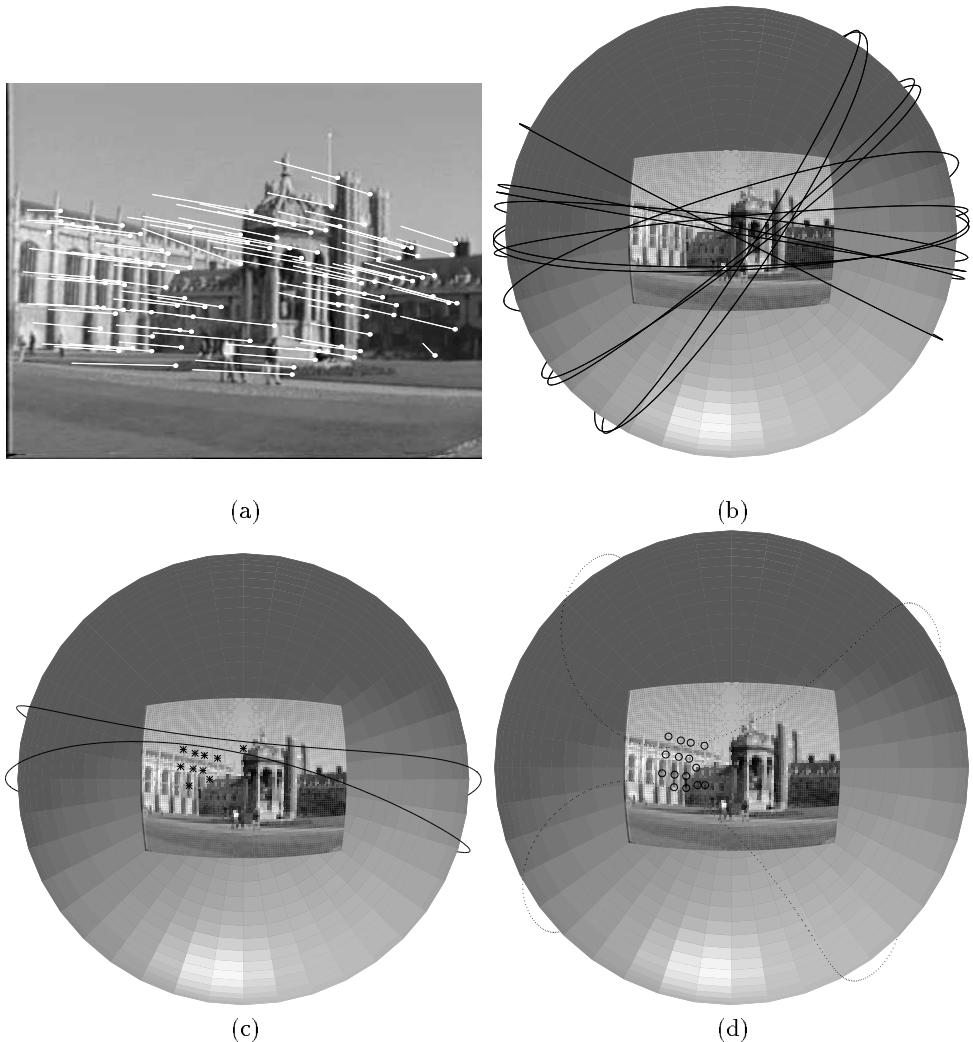


Fig. 1. Subspace constraints and uncertainty estimates: (a) shows the image feature displacements measured between two frames. The features were automatically extracted and matched and approximately 6 are incorrect correspondences. The translation of the camera was along the path just visible in the bottom right corner, but the image motion is mostly due to rotation. (b) shows the set of constraints obtained from 14 image regions by the Linearised Subspace Method. (c) shows the 95% uncertainty band and the image points contributing to the constraint. (d) shows an excessively uncertain constraint arising from distant points which are approximately on a plane

4 Constraint Combination

The problem of finding the epipole now becomes finding the direction or image point which is the best intersection of the set of uncertain constraints. This must be done after “rogue” measurement or outliers are rejected.

The uncertainty of the epipole estimate can be minimised by weighted least squares estimation. Generalising the constraints as unit vectors $\{\mathbf{n}_i\}$, we must minimize:

$$\sum_i w_i (\mathbf{e} \cdot \mathbf{n}_i)^2 = \mathbf{e}^T \left(\sum_i w_i \mathbf{n}_i \mathbf{n}_i^T \right) \mathbf{e} = \mathbf{e}^T \mathbf{A} \mathbf{e} \quad (13)$$

Assuming the statistical independence of the constraints, the weighting should be proportional to the inverse of the scalar variance [9].

An initial estimate of the epipole (eg. from an unweighted estimate) is needed so that the scalar weight can be determined only by the significant component of the uncertainty (parallel to the epipole). Since the constraint uncertainties are often highly anisotropic, many iterations are often necessary. Alternatively, an approximate weight can be found by averaging the uncertainty for all possible epipoles (for example, using the trace of the variance matrix).

An attractive feature of this form of epipole estimation is that the weighted outer product sum, \mathbf{A} ,

$$\mathbf{A} = \sum_i \mathbf{A}_i = \sum_i \frac{1}{\sigma_i^2} \mathbf{n}_i \mathbf{n}_i^T \quad (14)$$

(where $\sigma_i^2 = \mathbf{e}_0^T \mathcal{V}[\mathbf{n}_i] \mathbf{e}_0$ is the variance of the constraint normal in the (estimated) direction of the epipole, \mathbf{e}_0) encodes the estimate of the epipole, \mathbf{e} , and the variance (uncertainty) of the estimate, $\mathcal{V}[\mathbf{e}]$. The estimate of the epipole, \mathbf{e} , minimises $\mathbf{e}^T \mathbf{A} \mathbf{e}$ while $\mathcal{V}[\mathbf{e}]$ is the pseudo-inverse of \mathbf{A} . An example is shown in Figure 3 (c). Before being able to combine constraints it is important incorrect measurement(outliers) of visual motion due to failures in tracking are rejected.

5 Outlier Rejection

There are many reasons why some of the visual motion measurements will be incorrect and should not be used in the structure-from-motion calculation. Incorrect visual motion measurements arise from poor matching or tracking of features and independently moving objects. They can be removed by trying to form a solution from as many measurements as possible without using any that are not consistent with the solution found.

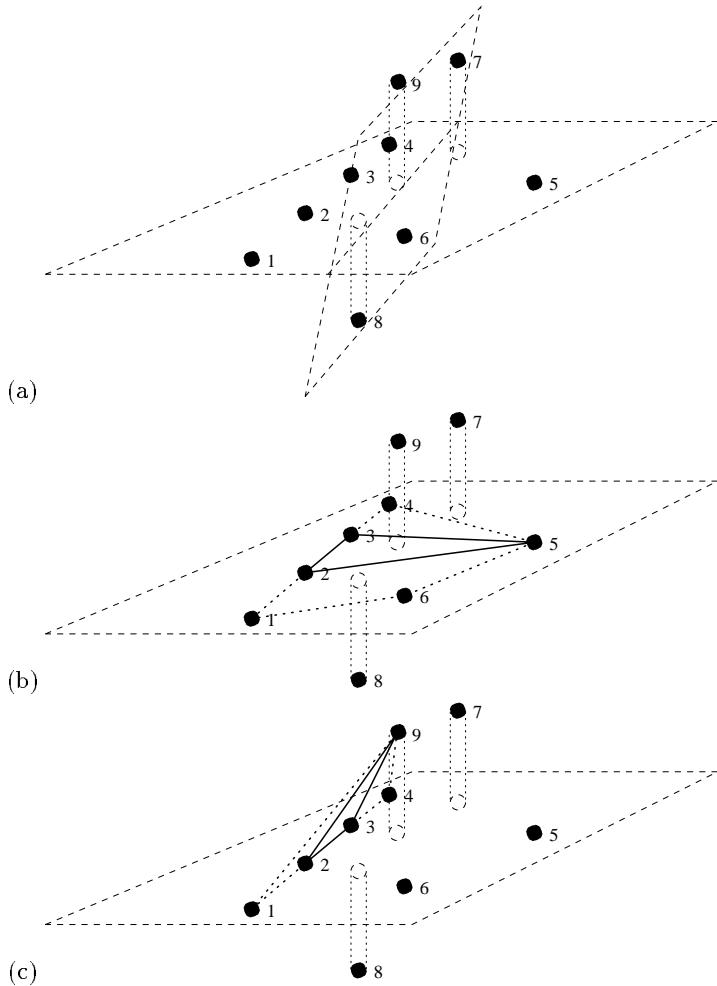


Fig. 2. Example rejection problem: Because the rejection algorithms considered in this section are generally applicable, a low-dimensional example problem will be used to introduce them: find a plane that fits a large number of the 3D points in a set. Each picture above shows the six inliers (1-6), the three outliers (7-9), and the plane. (a) **Poor initialisation using whole set:** Outliers are often of greater magnitude than true measurements, and can therefore greatly influence any initialisation using the whole set (as shown). In this case, take-one-out cross-validation would also fail because we have three significant outliers. (b) **Good solution from random initialisation set:** If the assembling of the plane estimate is initialised by three correct points (eg. 2, 3 and 5 as shown) then the other points can also be found. An average of four random trials (approximately) will be needed to find the plane. (c) **False solution from degenerate structure:** Solutions can also appear to be found when the verification comes from degenerate structure. Points 1 and 4 appear to be verifying the plane found from 2, 3 and 9, because they form a degenerate structure (a line) with all but one of these points. This effect should not cause trouble when a good solution can be found, as it will receive better verification (as in (a)). However, when the scene is truly degenerate, this can cause an (often strong) false positive.

5.1 Approaches to Rejection of Outliers

The aim of every rejection algorithm is to find a self-consistent set of “elements” with as few outliers as possible. However there are a large number of algorithms that can be used. Figure 2 presents an example problem of finding a plane of points in the presence of outliers, which will be used to demonstrate the strengths and weaknesses of a number of algorithms.

One method for rejecting outliers from any algorithm, is to start with the complete set of data, and then to reject the subset that does not *agree* sufficiently with the estimate formed from the set. This can be iterated until the set is stable. Though this seems sensible, large outliers can sufficiently bias the initial result so that they are never removed (see Figure 2a). A better approach, which was used in earlier papers on Affine Motion Parallax [10], is to disassemble the estimate using *take-one-out cross-validation* [15]. Each element is removed in turn and is checked against an estimate formed by the others. The constraint that agrees least is removed and the process repeated, until all constraints agree sufficiently. Though this performs better, the computational cost is high, and a few outliers can sufficiently bias the initial result so that removing one is still not sufficient (see Figure 2a).

A different method is to *assemble* the estimate, starting from an estimate from a minimal set, and combining only those constraints that agree. If too few constraints agree then a new initial estimate is needed. RANSAC [3] is a popular method that generalises this approach (see Figure 2b). Generally, this will be more efficient than the Hough methods. The problem here comes when the initial estimate is uncertain. If the uncertainty of the initial solution is great, then it is probably best to find a new initialisation, rather than build from this one. Even if the uncertainty is reasonable, there are probably still a number of erroneous elements that could be included, and different results may be found depending on how many of the agreeing elements are combined at once.

Each of the above methods can be iterated or two or more may be cascaded. Here we use a minimal basis to provide an initial candidate set free from gross outliers, which is then depleted by iteratively removing those elements that do not agree with a combined estimate.

5.2 Alternative Stages of Rejection

Most structure-from-motion algorithms that use outlier rejection, attempt to form a full projective egomotion solution and then use the epipolar constraint to remove outliers [14]. However, if the egomotion determination is broken into stages, as it has been in the algorithms presented here (finding the constraints and then the epipole before the egomotion), then the rejection can also be split into stages. Here, in addition to checking if the individual feature motions comply with the projective egomotion found for the set, we can also test if the individual features in small regions agree with the constraint found from that region, and if the constraints agree with one another. In this section we analyse and then test each stage in turn, in an attempt to determine if any advantage is gained from this approach.

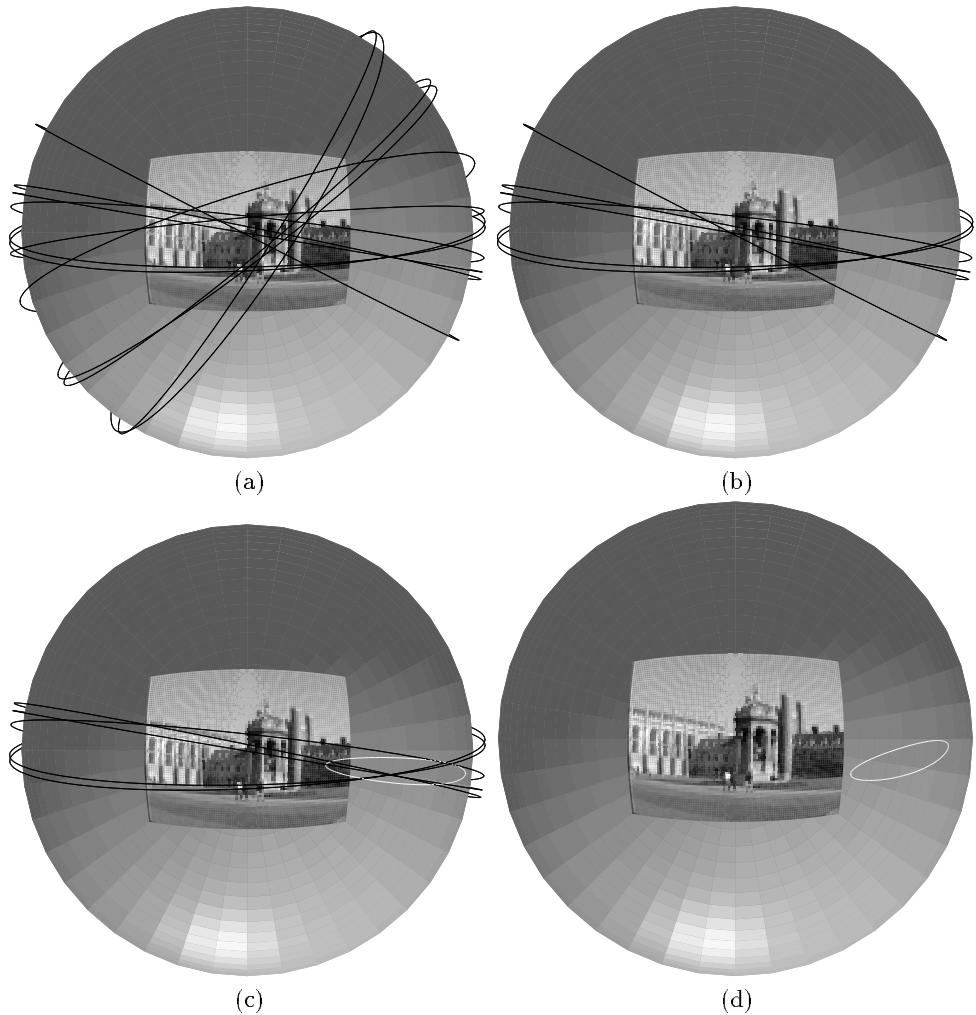


Fig. 3. The outlier rejection constraints: This figure gives a limited demonstration of the Consistency constraint and excessive error rejection for the features in Figure 1. (a) shows all the constraints found, (b) shows those with a reasonably small uncertainty and (c) shows all but one being accepted as agreeing with an epipole estimate, which is also shown as its 95% certainty ellipse. (d) shows the epipole ellipse given by the Fundamental Matrix calculated with the inliers and constrained to the epipole from (c).

The Affinity Constraint The Linearised Subspace method constraints (amongst others) are formed using the assumption that each comes from a region which was deforming affinely. This assumption constrains the feature motions – this can be seen more easily by considering the Affine Fundamental Matrix. It is therefore possible to test whether all of the features in the region comply with this assumption, and to reject those which do not. This could in fact be used as a perceptual grouping stage for any set of feature motions detected, regardless of the structure-from-motion algorithm that will use them, so long as the small motion approximation applies.

Though the algorithm initially seems simple to formulate, care must be taken to ensure that a few outliers are not masking a degenerate situation. Using the analogy of finding a plane of points in 3-space, the degenerate case is when there is a line of points plus outliers. Using RANSAC, we might select two points from the line and one outlier, and then validate this selection with the other points from the line (see Figure 2c). Though we could accept this erroneous set (in the hope of finding inconsistencies later) it is better to ensure that all planes found are contain at least four points with no three collinear.

The Consistency Constraint The second rejection stage that can be used is when the epipole constraints are combined: the constraint vectors should be coplanar if they are accurate, and the problem becomes finding the epipole estimate which is normal to the most constraint estimates. This is simpler than the Affinity Constraint, because degeneracy is more obvious, but is complicated by the tendency with all these constraint algorithms for correct and erroneous line constraints to pass through the image. This produces a bias in the solution but it also produces more candidate epipole estimates (intersections). For this reason, it is highly beneficial if the rejection methods discussed in the sections above and below are efficient. Another problem with this stage is that the relationship with the individual features has been lost.

The Epipolar Line Constraint Once the epipole has been estimated, then the projective “rotation” matrix can be found, and the rotational component of image motion predicted for each feature. The remaining image motion should be the translational component which will be towards (or away from) the epipole but have unknown magnitude (due to unknown feature depth). Any remainder not in this direction is evidence of an incompatibility. The effect should be similar to the Affinity Constraint above.

This is the rejection method most commonly used as the Fundamental Matrix finds the epipolar lines directly [14]. Here we must ensure that the Fundamental Matrix complies with the epipole found. It should therefore be formed from only the inliers from the previous stages, and should then be forced to have the same epipole. This is achieved by projecting the Fundamental matrix and the constraints on the Fundamental matrix into the subspace compatible with the epipole found. The matrix obtained can also be decomposed to give the camera motion estimate.

The Positive Depth Constraint The previous stages all effectively used the epipolar constraint, each time updating the epipolar line to include more evidence. Another useful though often ignored constraint is that all objects must be in front of the camera, and therefore all have translational motions in the same sense: either all towards or all away from the epipole. This can be applied in this scheme with the Epipolar Line Constraint. However, an egomotion estimate is needed to determine the rotational component of the visual motion, and therefore the Fundamental Matrix alone is not sufficient.

5.3 Using Uncertainty Estimates in Rejection

Uncertainty estimates can be used to provide intelligent bounds on the rejection or acceptance in the algorithms above: the χ^2 test on the Mahalanobis distance provides a threshold value, the confidence bound of which can be set [4]. This is the strictest use we have for the uncertainty estimates, as we are now accepting or rejecting on their absolute values, and not just weighting according to their relative values. If the uncertainty estimates are considered inaccurate then a “safety factor” can be used to reduce either false positives or negatives. However this does not detract from the advantages of having a threshold value that reflects the uncertainty in the element and the proposed solution.

6 Conclusions

Figures 4 and 5 show two contrasting situations. In the first, the epipole is approximate in the centre of the image. Most existing algorithms would have little difficulty in localising this epipole. There are a number of incorrect feature motion measurements, but both the (linear, non-iterative) Fundamental Matrix and the Linearised Subspace methods give good results. In the second however, the epipole is well right of the image. Here the Fundamental Matrix estimate is wrong, again placing the epipole in the image. The Linearised Subspace method gives an accurate result though, demonstrating our belief that epipole constraints can be both more accurate and more robust. The uncertainty predicted for the epipole estimate is high, but in both cases, applying this constraint to the Fundamental Matrix will ensure that the decomposition is correct. Another benefit of our approach is the uncertainty estimates are accurate.

Future work will concentrate on how best to select the small regions to give good coverage without excessive computation, and whether rigorous optimal combination is computationally viable. We conclude that the Fundamental Matrix benefits considerably from more direct estimates of the epipole. We have presented a novel epipole constraint algorithm and shown how to combine the constraints. We have also shown that outlier rejection is more reliable if it is performed at each stage of the calculation.

References

1. R. Cipolla, Y. Okamoto, and Y. Kuno. Robust structure from motion using motion

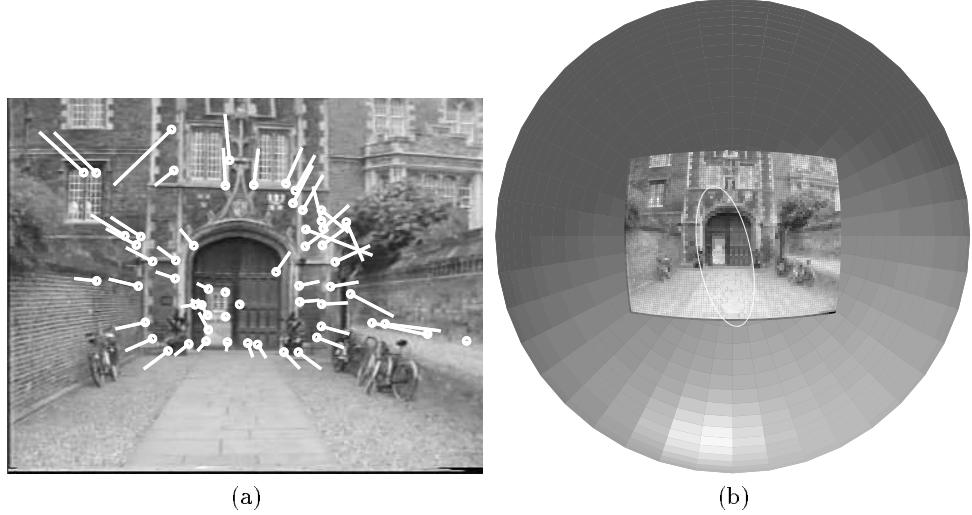


Fig. 4. Example needing considerable rejection: (a) shows the flow measured between two frames using approximate calibration. This motion, towards the centre of the doors with little rotation, is easier than the previous motion for most algorithms. However, a considerable proportion of features have incorrect correspondences. (b) shows the epipole estimated after using the Linearised Subspace method and the rejection methods outlined.

- parallax. In *Proc. 4th International Conf. on Computer Vision*, pages 374–382, 1993.
2. O.D. Faugeras, Q.-T. Luong, and S.J. Maybank. Camera self-calibration: Theory and experiments. In G. Sandini, editor, *Proc. 2nd European Conf. on Computer Vision*, pages 321–334. Springer-Verlag, 1992.
 3. M. A. Fischler and R. C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981.
 4. G.H. Golub and C.F. van Loan. *Matrix Computations*. North Oxford Academic, 1986.
 5. C.G. Harris. Geometry from visual motion. In A. Blake and A. Yuille, editors, *Active Vision*, chapter 16. MIT Press, 1992.
 6. R.I. Hartley and P. Sturm. Triangulation. In *Proceedings of the DARPA Image Understanding Workshop*, pages 957–966, 1994.
 7. A.D. Jepson and D.E. Heeger. Linear subspace methods for recovering translational direction. In L. Harris and M. Jenkin, editors, *Spatial vision in humans and robots*. Cambridge University Press, 1993. Also, University of Toronto, Department of Computer Science, Technical Report: RBCV-TR-92-40, April 1992.
 8. K. Kanatani. Computational projective geometry. *Computer Vision, Graphics and Image Processing (Image Understanding)*, 54(3):333–348, 1991.
 9. K. Kanatani. Renormalization for unbiased estimation. In *Proc. 4th International Conf. on Computer Vision*, pages 599–606, 1993.
 10. J.M. Lawn and R. Cipolla. Robust egomotion estimation from affine motion-

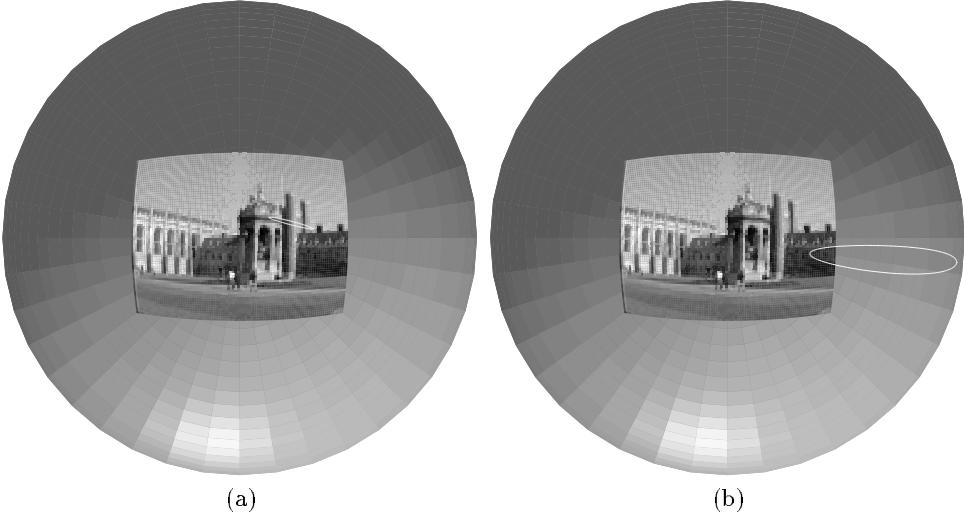


Fig. 5. Example with challenging motion: Fig. 1.(a) shows the flow measured between two frames. Half a dozen instances of poor feature correspondences are visible, mostly in the top left corner. The translation was along the path just visible in the bottom right corner, but the image motion is mostly due to rotation. (a) shows the epipole and uncertainty found by the linear Fundamental Matrix method and rejection, which is assigned a low uncertainty but is incorrect, and (b) the same using the Linearised Subspace method and the rejection methods outlined, which is very close to the actual epipole.

- parallax. In *Proc. 3rd European Conf. on Computer Vision*, pages I.205–210, 1994.
11. J.M. Lawn and R. Cipolla. Reliable extraction of the camera motion using constraints on the epipole. Technical Report CUED/F-INFENG/TR.240, University of Cambridge, 1995.
 12. H.C. Longuet-Higgins and K. Prazdny. The interpretation of a moving retinal image. *Proc. of the Royal Society of London, Series B*, 208:385–397, 1980.
 13. C. Rothwell, G. Csurka, and O. Faugeras. A comparison of projective reconstruction methods for pairs of views. In *Proc. 5th International Conf. on Computer Vision*, pages 932–937, 1995.
 14. P.H.S. Torr. *Motion Segmentation and Outlier Detection*. PhD thesis, Department of Engineering Science, University of Oxford, 1995.
 15. D.H. Wolpert. Combining generalizers using partitions of the learning set. In L.Nadel and D.Stein, editors, *1992 Lectures in Complex Systems*. Addison-Wesley, 1993.